

LOW LEAKAGE ASYMMETRIC SRAM CELL DEVICES

Cross-reference to Related Applications

[0001] This application claims the benefit of previously filed U.S. Provisional Patent Application Serial No. 60/554,198 filed on March 18, 2004 entitled, "LOW LEAKAGE ASYMMETRIC-CELL SRAM".

Field of the Invention

[0002] The present invention relates generally to SRAM (Static Random Access Memory) devices, and more particularly to low leakage power SRAM devices having device performance comparable to conventional SRAM devices.

Background

[0003] As a result of technology trends and the increased importance of portable electronic devices, leakage (static) power dissipation has emerged as a high priority design consideration in high-performance processor design. Historically, architectural innovations for improving performance relied on exploiting ever larger numbers of transistors operating at higher frequencies. To keep the higher resulting switching power dissipation at bay, successive technology generations have relied on reducing the supply voltage. In order to maintain performance, however, this has required a corresponding reduction in the transistor threshold voltage. Since the Metal Oxide Semiconductor Field Effect Transistor (MOSFET) sub-threshold leakage current increases exponentially with a reduced threshold voltage, leakage power dissipation has grown to be a significant fraction of overall chip power dissipation in modern, deep-submicron ($< 0.18 \mu\text{m}$) processes. Moreover, it is expected to grow by a factor of five every newer chip generation. For processors it is estimated that in $0.1 \mu\text{m}$ technology, leakage power will account for about 50% of the total chip power.

[0004] Since leakage power is proportional to the number of transistors, and given the projected large memory content of future System-on-Chip (SOC) devices, it becomes important to focus on Static Random Access Memory (SRAM) structures such as caches, which comprise the vast majority of on-chip transistors in some systems. Existing circuit-level leakage reduction techniques are oblivious to program behavior, such as how many bits to be stored will be high or low, and trade off performance for reduced leakage where possible. Combined circuit and architecture-level techniques reduce leakage for those parts of the on-chip caches that remain unused for long periods of time (for example, such as for thousands of cycles). The mechanisms that identify which cache parts will be unused and that enable leakage reduction incur considerable power and performance overheads that have to be amortized over long periods of time. As a result, these methods are not effective when most of the cache is actively used.

[0005] There is a need for SRAM storage with reduced leakage power while having comparable performance characteristics. As such, power consumption may be minimized while still providing the performance required in new generation systems and consumer devices.

Summary

[0006] The present invention seeks to satisfy at least some of the above unmet needs. Embodiments of the present invention include a family of improved asymmetric SRAM cell designs that can be used in new SRAM and cache memory designs referred to as the Asymmetric-Cell Caches (ACC). ACCs offer drastically reduced leakage power compared to conventional caches even when there are few parts of the cache that are left unused. ACCs exploit the fact that in ordinary programs most of the bits in caches are zeroes for both the data and instruction streams. It has been shown that this behavior persists for a variety of programs under different assumptions about cache sizes, organization and instruction set architectures, even when assuming perfect knowledge of which cache parts will be left unused for long periods of time.

[0007] Conventional SRAM cells are symmetrically composed of transistors with comparable leakage and threshold characteristics. The asymmetric SRAM cell designs of the present invention offer low leakage with little or no impact on latency. In asymmetric SRAM cells, selected transistors are "weakened" with respect to other transistors used in SRAM cells to reduce leakage power when the cell is storing a zero binary state (the most common case). Transistor weakening may be achieved by using higher voltage threshold (V_t) transistors, by varying transistor sizes, combinations of these approaches, or other means.

[0008] In addition to improved SRAM designs, the present invention also describes a novel sense amplifier (SA) design that exploits the asymmetric nature of our cells to offer cell read times that are comparable with conventional symmetric SRAM cells. Moreover, an embodiment of the present invention further presents a cache memory design based on ACCs that when compared to a conventional cache, the cache memory architecture of the present invention offers leakage reduction while maintaining high performance and comparable noise margins and stability.

[0009] In one embodiment of the present invention there is disclosed an asymmetric SRAM cell for storing a binary variable. The asymmetric SRAM cell exhibits reduced leakage power with respect to a comparable symmetric SRAM cell when the asymmetric SRAM cell stores a binary variable representing a predetermined binary value, such as a binary one or binary zero. The asymmetric SRAM cell is made up of a plurality of transistors of a first and second type operably coupled and configured as an asymmetric SRAM cell. At least one of the second type of transistor is made weaker than at least one of the first type of transistor. The two types of transistors are then variously configured such that the asymmetric SRAM cell achieves reduced leakage power with respect to a symmetric SRAM cell having the first type of transistor only.

[0010] The second type of transistor can be made weaker than the first type of transistor in various ways. One way is to increase the voltage threshold as compared to the voltage threshold of the first type of transistor. Another way is to decrease the channel width as compared to the channel width of the first type of transistor. Yet another way is to increase the channel length as compared to the channel length of the first type of transistor. Further, combinations of the above ways to make transistors relatively weaker, as well as other ways to make transistors relatively weaker may be used.

[0011] In another embodiment of the present invention there is disclosed a sense amplifier (SA) that exploits the characteristics of the asymmetric SRAM cell. A sense amplifier is coupled with

an asymmetric SRAM cell and provides faster access times when the asymmetric SRAM cell stores a first predetermined binary value. The sense amplifier is comprised of a first pair of cross coupled inverters across a bitline (BL) and a bitline bar (BLB) and a second pair of cross coupled inverters operably coupled with the first pair of cross coupled inverters. This is conventional up to this point. The present invention sense amplifier further includes a plurality of additional transistors forming a dummy column of cells that store a second predetermined binary value at all times wherein during a read operation of the SRAM cell one of the dummy cells will have its wordline asserted. The dummy column of cells are operably coupled with the first pair of cross coupled inverters. The sense amplifier is driven by four inputs operably coupled with a subset of transistors. The inputs include the BL and BLB that derive from the SRAM cell, as well as a dummy bit line (D), and a dummy bitline bar (DB). The D and DB are input to the dummy cells such that D is input to the sense amplifier on the same side as BLB while DB is input to the sense amplifier on the same side as BL.

[0012] Moreover, the transistors coupled with BL and BLB have higher transconductance characteristics than the transistors coupled with D and DB. This is achieved either by varying the threshold voltage or altering the size of the transistor channel widths or channel lengths.

[0013] In yet another embodiment of the present invention there is disclosed an SRAM device comprised of an array of SRAM cells wherein each SRAM cell stores a binary variable representing a predetermined binary value. In addition, each SRAM cell is an asymmetric SRAM cell having reduced leakage power with respect to a comparable symmetric SRAM cell as previously described. The SRAM device can be configured as a direct store SRAM device, a selectively inverted SRAM device, or a cache memory device. If the SRAM device is a cache memory device then it can either be configured as a direct store cache memory or a selectively inverted cache memory.

[0014] In still another embodiment of the present invention there is disclosed an asymmetric SRAM cell for storing a binary variable. The asymmetric SRAM cell exhibits reduced gate leakage power with respect to a comparable symmetric SRAM cell when the asymmetric SRAM cell stores a binary variable representing a predetermined binary value, such as a binary one or binary zero. The asymmetric SRAM cell is made up of a plurality of transistors of a first and second type operably coupled and configured as an asymmetric SRAM cell. At least one of the second type of transistor is made weaker than at least one of the first type of transistor. An additional pass transistor is included to reduce the voltage across the gate of a leaky transistor such that the asymmetric SRAM cell achieves reduced gate leakage power with respect to a symmetric SRAM cell.

Brief Description of the Drawings

[0015] FIGURE 1 illustrates an example circuit diagram of a conventional six transistor SRAM cell highlighting sources of leakage.

[0016] FIGURE 2 illustrates a circuit diagram of basic asymmetric SRAM cell, according to one embodiment of the present invention.

[0017] FIGURE 3 illustrates a circuit diagram of an asymmetric SRAM cell configured to address leakage, according to one embodiment of the present invention.

[0018] FIGURE 4 illustrates a circuit diagram of an asymmetric SRAM cell configured to address leakage, according to one embodiment of the present invention.

[0019] FIGURE 5 illustrates a circuit diagram of an asymmetric SRAM cell configured to address leakage and speed, according to one embodiment of the present invention.

[0020] FIGURE 6 illustrates a circuit diagram of an asymmetric SRAM cell configured to address leakage and speed, according to one embodiment of the present invention.

[0021] FIGURE 7 illustrates a circuit diagram of an asymmetric SRAM cell configured to address leakage and speed, according to one embodiment of the present invention.

[0022] FIGURE 8 illustrates a circuit diagram of an asymmetric SRAM cell termed a special precharge cell, according to one embodiment of the present invention.

[0023] FIGURE 9 illustrates a circuit diagram of an asymmetric SRAM cell termed a stability leakage enhanced cell, according to one embodiment of the present invention.

[0024] FIGURE 10 illustrates a circuit diagram of an asymmetric SRAM cell termed a stability speed enhanced cell, according to one embodiment of the present invention.

[0025] FIGURE 11 illustrates a circuit diagram of an asymmetric SRAM cell configured to address leakage through differences in transistor sizing, according to one embodiment of the present invention.

[0026] FIGURE 12 illustrates a circuit diagram of an asymmetric SRAM cell configured to address leakage and speed through differences in transistor sizing, according to one embodiment of the present invention.

[0027] FIGURE 13 illustrates a conventional sense amplifier.

[0028] FIGURE 14 illustrates a sense amplifier, according to one embodiment of the present invention.

[0029] FIGURE 15 illustrates a data flow diagram illustrating using selective inversion of byte data to optimize use of asymmetric SRAM cells, according to one embodiment of the present invention.

[0030] FIGURE 16 illustrates a circuit diagram of an asymmetric SRAM pass cell designed to reduce gate leakage, according to one embodiment of the present invention.

[0031] FIGURE 17 illustrates a portion of the pass cell when holding a logic "0".

[0032] FIGURE 18 illustrates a portion of the pass cell when holding a logic "1".

Detailed Description

[0033] Ideally, an SRAM cell should be fast and should dissipate low leakage power. This is increasingly at odds with the fundamental technology trade off between transistor speed and leakage. Conventional high performance SRAM cells use a symmetric configuration of six transistors with comparable threshold voltages. One can reduce leakage by using higher V_t transistors, but unfortunately using an all high V_t transistor cell degrades performance by an unacceptable margin.

[0034] The goal of the asymmetric SRAM cells of the present invention is to reduce leakage while maintaining high performance based on the following approach: select a preferred state and weaken only those transistors necessary to drastically reduce leakage when the cell is in that state. These cells exhibit asymmetric leakage and access behavior. Fortunately, their asymmetric access behavior can be exploited to maintain high performance while reducing leakage.

[0035] For purposes of illustration, the following convention will be used. A high V_t (HV) transistor is obtained from a basic 0.13 μm , 1.2V, transistor (referred to herein as the regular V_t (RV) transistor) by artificially increasing the V_t by 0.2V. 0.2V was chosen because it leads to a difference of about 10 times between the leakage currents of HV and RV transistors, which is typical of dual V_t technology. Those of ordinary skill in the art will realize that other relative changes to V_t can be implemented. The data values illustrated herein are but one example chosen to illustrate the results of the present invention when the asymmetric concept is applied. For illustration purposes, A "high V_t transistor" as used herein is defined as a transistor having a relatively higher " V_t " or threshold voltage than other transistors typically used in an SRAM cell design. The reason for selecting transistors having a higher V_t than others within the SRAM cell is to reduce the leakage current and thereby reduce an SRAM cell's leakage power. Although the high V_t transistor example described herein has a threshold voltage (V_t) which is 0.2volts higher, this is only an example for a 1.2volt, basic, 0.13 micron transistor. Different shifts of V_t could be used, using either higher or lower V_t differential voltages, so long as the leakage current draw is reduced as required for a given SRAM cell design or application. Additionally, transistors in technologies other than the basic 0.13 micron example can be used.

[0036] Moreover, the present invention has been described and illustrated using MOSFET type transistors. Those of ordinary skill in the art can appreciate that other types of transistors and the like can be substituted for MOSFETs.

[0037] FIGURE 1 illustrates a conventional SRAM cell 10 comprised of two inverters 12, 14, (P2, N2) and (P1, N1), and two pass transistors 16, 18, N3 and N4. In the inactive state, a wordline (WL) 20 is held low so that the two pass transistors 16, 18 are off isolating the cell from a bitline (BL) 22 and bitline-bar (BLB) 24. At this stage the bitlines 22, 24 are also typically charged at VDD (e.g., logic '1'). Cells spend most of their time in the inactive state. In this state, most of the leakage is dissipated by the transistors that are off and that have a voltage differential across their drain and source. The value stored in the cell (i.e., the cell state) determines which transistors these are. When the cell is storing a '0', as in FIGURE 1, the leaky (subthreshold leakage) transistors are P2, N1 and N3. If the cell were storing a '1' then transistors P1, N4 and N2 would dissipate leakage power. On gate direct tunneling leakage occurs in transistors P1 and N2 while edge directed tunneling leakage occurs in transistors P2, N3, N1, and N4. A simple technique for reducing leakage power would be to replace all transistors with high- V_t ones, but this unacceptably degrades the bitlines discharge times by 61.6%.

[0038] Since ordinary programs exhibit a strong bias in cache-resident bit values, another possibility to reduce leakage power, but at the same time keep read access times short, is to choose a preferred stored value and to only replace those transistors that contribute to the leakage power in this state with HV transistors. This is illustrated in FIGURE 2 where P1, N4 and N2 have been made weaker with respect to P2, N1, and N3. This basic asymmetric SRAM cell 25 was simulated and exhibits the same leakage as the RV cell 10 of FIGURE 1 when holding a logic '1', but its leakage is reduced by 70X when holding a logic '0.'

[0039] The read access time of the basic asymmetric cell is, however, degraded. Due to N2's and N4's higher threshold voltage, the bitline discharge takes longer. The discharge times for BLB and BL are 12.2% and 46.4% longer than the discharge time for the RV cell, respectively. Discharge time is defined as the time from when the wordline is raised to when one of the bitlines reduces to 90% of its precharge value. The number 90% was chosen due to it being an appropriate differential signal for sense amplifiers to trigger.

[0040] P-Channel Metal Oxide Semiconductor (PMOS) transistors have very little effect on a cell's read access time because the role of pulling down the bitlines is played by the two n-channel Metal Oxide Semiconductor (NMOS) transistors on the side of the cell storing the '0'. Thus, a better asymmetric cell can be configured using the basic asymmetric cell of FIGURE 2 with P2 also set to high V_t . This cell, shown in FIGURE 3, is referred to as the Leakage Improved 2 (LI2) cell 30 and has the advantage of partially reduced leakage in the high leakage state. When the cell is holding a logic '1' its leakage is reduced by 1.6X relative to the RV cell, and when holding a logic '0' its leakage is reduced by 70X. The discharge times for BLB and BL are 12.2% and 46.4% longer than the discharge times for the RV cell, respectively, the same as the basic asymmetric cell's discharge times.

[0041] A further improvement is possible since by using a sense amplifier (described below) that matches the read time on the slow side of the cell to the fast side, there is no need for N1 to be low V_t . This leads to the cell in FIGURE 4, referred to as the Leakage Improved 3 (LI3) cell 40 or leakage enhanced (LE) cell 40. This cell further reduces leakage in the high leakage state, so that its leakage relative to the RV cell 10 is reduced by 7X in the '1' state and by 70X in the '0' state. The BL discharge time is now 61.6% longer than the discharge time for the RV cell 10, but that is of minor importance due to the novel sense amplifier design, as we will see later. The two asymmetric cells, LI2 30 and LI3 40, take the basic asymmetric cell 25 of FIGURE 2 and improve its leakage performance without affecting its read access time.

[0042] Another design challenge is to take the basic asymmetric cell 25 and improve its read access time while keeping some of the leakage benefits of the basic asymmetric cell 25. To eliminate the speed penalty incurred in the basic asymmetric cell 25 due to both pull-down paths having one high V_t transistor, both N2 and N3 are kept at low V_t while P1 is made high V_t . This cell is shown in FIGURE 5 and is termed the Speed Improved 1 (SI1) cell 50. The SI1 cell 50 has discharge times for BLB and BL which are 0% and 46.7% respectively longer than the RV cell 10. Thus one side of the cell is just as fast as the RV cell 10. However, this cell suffers from higher leakage than the basic asymmetric cell 25, with a leakage reduction of 2X relative to RV cell 10 when holding a '0', and no leakage reduction when holding a '1'.

[0043] The same transformations performed on the basic asymmetric cell 25 to improve its leakage performance can also be performed on the SI1 cell 50. First, P2 is made high V_t (FIGURE 6), and then N1 is also made high V_t (FIGURE 7). These two new cells are named Speed Improved 2 (SI2) 60 and Speed Improved 3 (SI3) 70, respectively. The SI2 cell 60 has leakage reductions of 2X and 1.6X when storing a '0' and '1', respectively, while the SI3 cell 70 has leakage reductions of 2X and 7X. The SI3 cell 70 is also referred to as the Speed Enhanced (SE) cell 70.

[0044] These two cells have no read access time degradation compared to the RV cell 10 along BLB, but have a 46.5% and 61.6% degradation along BL respectively. Once again, the degradation along BL is of minor importance due to the novel sense amplifier.

[0045] Note that the SE cell 70 reverses the preferred leakage state to the state when the cell is holding a '1'. All further references to this cell will have the '1' state as the preferred state so that the cell language remains in conformity with other cells. It should be noted that in practice the cell bitlines can be flipped to allow for '0' to be the preferred state without affecting any of the performance or stability results shown here.

[0046] One would like to combine the low leakage of the LI2 30 and LE 40 cells with a very small read access delay. Yet another asymmetric cell addresses these objectives, but it requires a different read operation. In the steady state, instead of keeping BL precharged to VDD, it is kept at ground. Now, N4 18 can be kept low V_t for the preferred '0' state. This is termed the Special Precharge (SP) cell 80 and it is shown in FIGURE 8. This asymmetric cell requires changes to the peripheral circuits of the SRAM array. Nevertheless, the results for this cell indicate that leakage is reduced by 83.3X in the '0' state, while the '1' state shows no leakage reduction. Bitline discharge times are degraded by 12.2% and 0%, respectively, for this example.

[0047] Until now, only the bitline discharge times of the different cells have been compared, and write times have been ignored. The write times of the cells are less important because stronger write drivers can be designed to drive the bitlines, and write drivers are a small portion of the total SRAM. The write times of the asymmetric cells all lie within the write times of the RV cell and the HV cell.

[0048] The LE cell 40 and SE cell 70 are the two best designs from the two sets of asymmetric cells as indicated by test results. Therefore, only these two cells, and variations on them, will be referenced in the remainder of this description.

[0049] Another major consideration with the cell design is its stability. There are two interrelated issues: read stability and noise margins. Read stability indicates how likely it is to invert the cell's stored value when it is being accessed. This is computed as the ratio of I_{trip}/I_{read} , where I_{trip} is the current through the pull-down NMOS when the state of the cell is being reversed by injecting an external current I_{test} , and where I_{read} is the maximum current through the pass transistor during a read.

[0050] The static noise margin (SNM) of an SRAM cell is defined as the minimum DC noise voltage necessary to flip the state of the cell. For the present invention, the stability of all cells was measured by simulation via both the Static Noise Margin (SNM) and the I_{trip}/I_{read} methods. Under both stability tests, the stability was first measured under nominal conditions, assuming no process variations. Then, to measure stability under process variations, two sets of tests were performed. First, the SNM and I_{trip}/I_{read} tests were performed on 59,049 combinations of different V_t and length variations for all six transistors in the cell. The combinations included modifying by $\{-3\sigma, 0, 3\sigma\}$ the NMOS transistors' V_t and length values and the PMOS transistors' V_t value. The worst case value for various cells was found, and compared to the worst-case value obtained for the RV cell.

[0051] Second, Monte-Carlo analysis was performed to obtain a distribution for the SNM and Itrip/Iread. For each cell, 500 scenarios for V_t and channel length were randomly generated, consistent with their joint distributions, and simulated. The mean of the distribution was estimated using the unbiased estimator in (1), and the variance was estimated by using the unbiased estimator in (2). Furthermore, the Normal Scores Method was used to graphically determine the distribution type. Given the distribution type, mean, and variance, the probability of failure for various cells was then computed.

[0052] The SNM of the LE 40 and SE 70 cells were computed through simulation. The SNM of the RV cell 10 was also computed to be used as a reference. Under nominal conditions, the SNM of the LE 40 and SE 70 cells were 0.246V and 0.221V, respectively, while the SNM of the RV cell 10 was 0.250V. Thus, the LE cell 40 and SE cell 70 show a decrease in SNM of 1.6% and 11.7%. One would expect that by using higher threshold voltage transistors in the design, the SNM of the cells would increase, but the asymmetry of the cells skews the lobes of the butterfly curve and decreases the SNM, as will be explained below.

[0053] First, let us examine the SNM of the cells when the wordline is not active. During this state, the SRAM cell is not as vulnerable as when it is being read, but a study of this case helps to understand the decrease in the SNM when the cell is being read. When the wordline is off, the only transistors that affect the SNM are the four transistors comprising the back-to-back inverters.

[0054] Since the four internal transistors of the LE cell 40 are all high V_t , the cell has equal low and high noise margins of 0.685V, a 22.6% increase over the standby SNM of the RV cell, 0.559V. However, when the SNM of the cell is being measured during a read the cell has high SNM in one state, 0.363V, and low SNM in the other, 0.246V. The asymmetry in the LE butterfly curve is due to the mismatch between the strength of the pass-gate (N3) and pull-down (N2) transistors. During a read, the N3 pass transistor 16, due to it being low V_t , has a higher conductivity than N2 and raises the voltage at the storage node to a higher voltage than if the two NMOS were of equal strength.

[0055] For the SE cell 70, the internal inverter pair are different. Thus the standby (i.e., with the wordline off) SNM of the cell has asymmetric lobes with noise margins of 0.535V and 0.727V, in the worst case a 4.2% decrease in noise margin compared to the RV cell. The source of this mismatch is the V_t difference between N1 and N2, which causes one of the transfer characteristics to commence its transition in the SNM plot from '0' to '1' later than normal. During a read, the mismatch between the size of the lobes becomes exaggerated because it is as if a constant is subtracted from the noise margin on each side of the cell since each side of the cell has equal strength pass transistors and pull-down transistors. While being read, the SE cell 70 has low and high noise margins of 0.222V and 0.365V respectively.

[0056] The asymmetric cells' stability performance degrades compared to that of the RV cell. Since process variations induce an asymmetry in the butterfly curve, the original asymmetry inherent in the butterfly curves for the LE 40 and SE 70 cells allows one lobe of the butterfly curve to become pinched off even further and lose stability. For the LE cell 40 the butterfly curve becomes pinched off when N3 becomes stronger than N2 and P1 increases in strength, while N1 does not. The

worst case for the SE cell 70 occurs at a different process corner. The butterfly curve becomes pinched off when P2 decreases in strength and N2 increases in strength, and N4 gets stronger than N1.

[0057] Monte-Carlo Analysis was also performed on the RV 10, LE 40 and SE 70 cells. The Normal Scores method reveals that the distributions for all cells were Gaussian. Due to their very small standard deviation, the SNM of all cells remains very close to their respective mean average. Thus the mean of the SNM becomes a very important measure, and is a better reflection of the stability than the nominal or worst-case SNM. Using the mean as a measure of stability, the LE cell 40 has a 7% increase in SNM and the SE cell 70 has a 5.8% decrease.

[0058] Using the SNM as a measure of stability showed that the LE cell 40 was comparable to the RV cell 10 while the SE cell 70 showed a marginal decrease in stability. When Itrip/Iread is computed by simulation, it is seen that the SE cell 70 outperforms the RV cell 10 and the LE cell 40 suffers.

[0059] The LE cell 40 has a lower Itrip/Iread value due to the Vt mismatch between the pass transistor and pull-down transistor on one side of the cell. The Itrip values from both sides of the cell show a drop compared to the Itrip value from the RV cell 10 due to both pull-down transistors becoming high Vt. However, with N3 16 remaining low Vt, Iread on the fast side of the cell does not suffer the same drop, and Itrip/Iread falls compared to that of the RV cell 10.

[0060] The SE cell 70, due to it having the same strength pull-down and pass transistors 16, 18 on each side of the cell, does not experience the same problem as the LE cell 40. On the slow side of the cell, both Itrip and Iread fall compared to the RV cell 10, but Iread falls by a larger amount thus increasing the Itrip/Iread. On the fast side of the cell, Iread does not change compared to the RV cell 10, but Itrip increases slightly. In the RV cell 10, the reduction in voltage (due to leakage) at the stored '1' node degrades the current sinking capacity of the pull-down NMOS. In the SE cell 70, because of the high Vt transistors on the '1' side of the cell there is no degradation in the current sinking capacity of the pull-down transistor and thus Itrip increases leading to a larger Itrip/Iread.

[0061] A total of 59,049 different corner cases of process variations were simulated and the worst case Itrip/Iread was noted in each cell. The LE cell 40 and the RV cell 10 achieve their worst-case Itrip/Iread for the same process corner: when the difference in strength between N2 and N3 is amplified with N2 becoming weaker, and N3 16 becoming stronger. The SE cell 70, however, suffers its worst-case Itrip/Iread when N4 18 becomes stronger than N1.

[0062] Monte-Carlo analysis show that Itrip/Iread is also Gaussian from the linear plots obtained from the Normal Scores Method. The standard deviation is very small and most cells will be very near the mean where the LE shows a 4.35% decrease and the SE cell 70 shows a 14.84% increase in Itrip/Iread.

[0063] The SE 70 and LE 40 cells have either a lower stability in the SNM test or the Itrip/Iread test. In many cases, the stability of the cell is a critical factor to obtain a desired yield and to lower the cost of the chip. In that regard, two derivative cells, one from the LE cell 40 and one from the SE cell 70, have been developed that improve upon their SNM, but do not decrease the leakage as much as the SE 70 and LE 40 cells. The two new cells are named Stability-Leakage Enhanced (SLE) 90 and Stability-Speed Enhanced (SSE) 100 and are illustrated in FIGURES 9 and 10 respectively.

[0064] One way to improve the SNM of the cells under process variations is to try to make the size of the lobes of the butterfly curve symmetric. For the LE cell 40 the lobes can be made more symmetric by making N2 low V_t , but this new cell would just be the SE cell 70. Another option is to make P1 low V_t . This change, shown in FIGURE 9, makes the lobes of the butterfly curve more symmetric. The SNMs are now 0.360V and 0.283V instead of 0.363V and 0.246V. To make the SE cell's 70 SNM plot more symmetric, P2 can be made low V_t yielding SNMs of 0.256V and 0.362V instead of 0.222V and 0.366V.

[0065] For these stability improved cells, all the previous tests for leakage, performance, and stability can be performed to compare them to the cells they were derived from, as well as to the RV cell 10.

[0066] The leakage performance of the stability improved SLE 90 and SSE 100 cells falls off, as expected due to one transistor in the LE 40 and SE 70 cells being re-converted to a low V_t transistor. For the SLE cell 90, the leakage reduction when holding a '1' remains unchanged at a 6.96X reduction relative to RV cell 10, but the leakage reduction when holding a '0' changes from 69.5X to 2.5X. For the SSE cell 100, when it is holding a '0' the leakage reduction stays at 2.04X, but when it is holding a '1' the leakage reduction changes from 6.96X to 1.91X.

[0067] Since the PMOS transistors do not play a large role in discharging the bitlines, it would be expected that the discharge time for the stability improved cells to be very close to the cells they derived from. Through simulation, it is seen that the discharge times along BL and BLB remain almost constant. As for the write times, SLE cell's 90 write time decreases to a 33.15% increase over RV cell's 10 write time from LE cell's 40 35.95% increase. The SSE cell's 100 write time jumps to a 49.22% increase over the RV cell's 10 write times.

[0068] A stability analysis has also been performed on the derivative cells for both the SNM and Itrip/Iread. Both derivative cells perform better than the RV cell 10 in the worst case, and under Monte-Carlo analysis. Under the Itrip/Iread method, there is very little change, because Itrip/Iread depends strongly on the NMOS transistors, which have not been changed, but the stability-improved cells perform slightly worse than the cells from which they were derived.

[0069] It has been shown that when stability is recovered through a change in threshold voltage of the PMOS transistors, a large portion of the leakage benefits of the asymmetric cells are lost. Furthermore, the Itrip/Iread of the LE cell 40 could not be improved by threshold voltage assignment. Another way of improving stability is to resize some of the transistors to reclaim the conductance lost due to the high V_t assignment. This change does not have a large effect on the leakage characteristics because leakage increases exponentially with reduced threshold voltages, but increases only linearly with transistor size. Moreover, the low Itrip/Iread of the LE cell 40 can be improved by transistor resizing.

[0070] The lobes of the SNM plot for the SE cell 70 can be made more symmetric by making N1 wider. In our case, we increased the width of this transistor by 26%, leading to a new cell shown in FIGURE 11 and referred to as Resized Speed Enhanced (RSE) 110. The SNM for the RSE cell 110 is comparable to that of the RV cell 10 and the change in N1's size leads to an increase of only 2.9% in cell area. The SNM margins are now 0.253V and 0.347V instead of 0.222V and 0.366V.

The RSE cell's 110 nominal value for Itrip/Iread does not change much compared to the nominal value for the SE cell 70. On the slow side of the cell, which had the higher Itrip/Iread value for the SE cell 70, the increase in N1's size allows for Itrip to become larger and increases the Itrip/Iread value. The fast side of the cell however, which has the limiting Itrip/Iread value, has a reduced Itrip that reduces the final value of Itrip/Iread to 2.53. The reduction in Itrip is due to the '1' storage node having a slightly lower voltage due to the increased leakage through N1. Nevertheless, the RSE cell's 110 Itrip/Iread value is still 11.8% better than that of the RV cell 10.

[0071] For the LE cell 40, increasing the width of N2 allows the conductance of N2 to approach that of N3 16, which leads to an increase in Itrip, thus increasing Itrip/Iread. By increasing N2's width by 22%, (leading to an only 2.4% increase in cell area) the Itrip/Iread value of the new Resized Leakage Enhanced (RLE) cell 120 (FIGURE 12) was made to be 2.28, which is comparable to the Itrip/Iread value of 2.26 of the RV cell 10. The increase in N2's width also increases the SNM of the RLE cell 120 where the margins are now 0.349V and 0.280V instead of 0.363V and 0.246V.

[0072] As expected, the leakage performance of the resized cells is better than that of the SLE 90 and SSE 100 cells. For the RLE cell 120 the leakage reduction when holding a '1' remains unchanged at a 6.96X reduction relative to RV cell 10, but the leakage reduction when holding a '0' only slightly reduces from 69.5X to 57.9X. The SLE cell's 90 leakage reduction when holding a '0' was only 2.5X. When the RSE cell 110 is holding a '0' the leakage reduction stays at 2.04X relative to RV cell 10, and when it is holding a '1' the leakage reduction only changes from 6.96X to 6.79X. This change is also minimal when compared to the SSE cell's 100 leakage reduction of 1.91X.

[0073] Due to the increased size of the pull-down NMOS transistors, the resized cells have the potential of improving the read-access time of the cell. For the RLE cell 120 the discharge time along BLB remains at a 61.1 % increase over the RV cell's 10 BLB discharge time, but the BL discharge time is now only 3.7% longer than the RV cell's 10 discharge time. As noted previously, only the BL discharge time is important due to the timed read based on a new sense amplifier. For the RSE cell 110, the discharge time along the fast side of the cell, BL, does not change, but the discharge time along BLB is reduced from the SE cell's 70 61.7% increase over RV cell 10 to a 49.2% increase over RV cell 10. This extra performance along BLE plays no important role in the cell's performance. As for the write times, the RLE cell's 120 write time increases to a 39% increase over RV cell's 10 write time from LE cell's 40 35.95% increase. The RSE cell's 110 write time jumps to a 45% increase over RV cell's 10 write times.

[0074] The stability analysis has also been performed on the resized cells for both the SNM test and Itrip/Iread test. Both resized cells perform better than the RV cell in the worst case, and under Monte-Carlo analysis for the SNM. Under the Itrip/Iread test, the RLE cell 120 now performs better than RV cell 10 both in the worst-case and on average. The increase in N1's size accomplishes the higher Itrip/Iread. The RSE cell's 110 Itrip/Iread value also increases slightly under all tests, even surpassing the SE cell's 70 Itrip/Iread value in the worst case. With a larger pull-down transistor, the process variations do not have as much an effect on the RSE cell's 110 stability.

[0075] Another figure of merit for the different cells is their stability under different supply voltages. For the technology being used, the nominal supply voltage is 1.2V. Monte-Carlo analysis

has been performed for the RV 10, LE 40, SLE 90, RLE 120, SE 70, SSE 100 and RSE 110 cells for supply voltages ranging from 0.75V to 1.6V.

[0076] For voltages above 1.2V, LE 40, SLE 90 and RLE 120 improve their SNM advantage over the RV cell 10. With a higher VGS, the difference in conductance between the pass-gate (N3) and pull-down (N2) transistors, which was the root cause of the low stability at 1.2V, diminishes. At higher voltages, the SNM of the SE 70 and SSE 100 cells starts to diminish just as the SNM of the RV cell 10 but at a lower rate. The SNM of the RSE cell 110 levels off at higher voltages.

[0077] With lower supply voltages, the SNM of the asymmetric cells starts to suffer. For the LE 40, SLE 90 and RLE 120 cells, the SNM decreases rapidly, but the SLE cell's 100 SNM remains comparable to that of RV cell 10, while the RLE cell's 120 SNM becomes comparable to that of the LE cell's 40. This decrease in stability is caused by the difference in conductance between regular voltage and higher voltage transistors at low VGS's. Furthermore, at low VGS, the extra conductance of the larger transistor in the RLE cell 120 does not have a large effect since the transistor is not fully on. The SNM of SE 70, SSE 100 and RSE 110 also decreases, but not as fast as that of the LE cell 40. Again, this decrease in SNM is due to the difference in conductance at low VGS's.

[0078] The same tests were performed for the Itrip/Iread method with the result that the curves for all cells are much better behaved. The SE 70 and SSE 100 cells have a near 24% advantage over the RV cell at 0.75V and an 8% advantage at 1.65V. The LE 40 and SLE 90 cells have approximately a 16% decrease in Itrip/Iread at 0.75V and are comparable at 1.65V to the RV cell 10. The resized cells behave slightly differently, with the RSE cell 110 having an 11.7% improvement at 1.65V and a 32.2% improvement at 0.75V. The RLE cell 120 has a 9.6% improvement at 1.65V and a 4% decrease at 0.75V.

[0079] A conventional sense amplifier 130 is shown in FIGURE 13. It is not suitable for the present invention due to the slow access time when the cell is storing a '0'. To obtain fast read times regardless of the data value, a new sense amplifier 140 has been designed and is shown in FIGURE 14. Compared to the conventional sense amplifier 130, the new sense amplifier 140 has four additional transistors 142, 144, 146, 148 and an area increase of roughly 0.229 μm^2 or 14.4%.

[0080] In addition to BL 132 and BLB 134, the sense amplifier 140 has two new inputs, D 150 and DB 152. These are connected to a dummy column of cells that store '1' at all time, but which are otherwise exactly identical to all other cells in the array. This dummy column extends the full length of the SRAM array such that during every read operation, one of the dummy cells will have its wordline asserted. Since the dummy cells always store a '1', they are always fast on the discharge (as fast as the fast side of any other cell), and they are used to provide something like a timer signal. This is achieved by connecting the dummy bitlines 150, 152 to the sense amplifier 140 in a reverse way. D 150 is connected to the right side, where BLB 134 is connected, and DB 152 is connected to the left side, where BL 132 is connected. This enables D 150 and DB 152 to trigger a fast read of a '0' result when the cell being read has a '0' content.

[0081] Sensing a '1' is as fast as a conventional sense amplifier 130 since this is done by sensing a discharge of BLB 134 due to the action of the fast side of the cell. Sensing a '0' is initiated at a later time than it would be in a conventional sense amplifier 130 to allow sufficient time for the fast

side to trigger the sense amplifier 140 if it has to do so. While initiating the sensing for a '0' is delayed, the combined effect of the dummy cell and the slow side of the asymmetric cell makes the sensing process itself much faster once initiated, so that the end result becomes available at about the same time as it would when sensing a '1'.

[0082] The detailed operation of the sense amplifier 140 is as follows. Initially, the bitlines 132, 134 are precharged and all four amplifier inputs rise to VDD. During this phase the sense amplifier 140 is being reset and nodes A and B are reset to an intermediate value. During a read operation, either BLB 134 will discharge (cell has a '1', fast discharge from the fast side) or BL 132 will discharge (cell has a '0', slow discharge from the slow side). Furthermore the signal DB 152, which is on the fast side of the dummy cell, will be discharged since the dummy cells permanently hold a logic '1'. If BLB 134 is being discharged a logic '1' is being sensed and the differential pair comprised of N1 and N2 causes increased current to pass through the left branch, thus increasing the voltage at node B and decreasing the voltage at node A. Through the positive feedback loop of P1, P2, N5, and N6, the rate of change for nodes A and B are increased to achieve quick sensing. When BL 132 is being discharged a logic '0' is being sensed. It does so at a slower rate since it is being discharged from the slow side of the asymmetric cell. To achieve fast sensing in this case, the dummy bitlines 150, 152, which are connected to the differential pair of N3 and N4, initiate the sensing of a logic '0'. Through the combined effect of DB 152 and BL 132 being discharged, albeit at a slower rate, approximately symmetric sense times are achieved.

[0083] For this sensing scheme to achieve reliable results it must allow for adequate time for BLB 134 to discharge before initiating a logic '0' read. This safety factor is achieved in two ways. First, the dummy bitlines 150, 152 are connected to all sense amplifiers and therefore have a slightly higher capacitive load compared to real bitlines 132, 134 leading to a slower discharge on DB 152 compared to BLB 134. The extra capacitive loading does not slow the sense time when BL 132 is discharging because of the concerted effort between BL 132 and DB 152 to sense the same value. Second, the transistors connected to the bitlines 132, 134 are wider than the transistors connected to the dummy bitlines 150, 152 leading to a higher transconductance and higher gain from the bitlines 132, 134 to the output than from the dummy bitlines 150, 152.

[0084] To limit the sense power, the sense amplifiers are clocked. The sense clock turns on the amplifiers and sets them up in their high gain region before the sensing occurs. To improve yield and ensure low-power operation, the clock path is matched to the data path. Matching is achieved by using an extra set of dummy bitlines to match the bitline delay and clock the sense amplifiers at the appropriate time.

[0085] Using the above cells and the sense amplifier 140 presented above, a 32-Kbyte SRAM example was designed and simulated to measure leakage, and read and write times. Each of the 128 SRAM sub-arrays contains 64 cells along each bitline, and 32 cells along each wordline. The SRAM was simulated at a temperature of 110° C with the RV cell 10, basic asymmetric, LE 40, SLE 90, RLE 120, SE 70, SSE 100, RSE 110 and HV 25 cells. Furthermore, the RV 10 and HV cells 25 were simulated with a conventional sense amplifier 130, and these results were used as a reference for our design.

[0086] The leakage trends seen above for the single cell remain true for the complete SRAM, where the LE 40 and SE 70 cells offer a reduction of 70X and 2X while storing a '0' and a reduction of about 7X when storing a '1.' The stability improved cells, and the resized cells also show the same leakage trends from the single cell experiments.

[0087] The total SRAM read access time includes four components: 1) input register propagation delay and hold times; 2) the address decoding delay; 3) the delay for wordline, bitline and sensing; and 4) the output register setup time. Only the delay for wordline, bitline and sensing is affected by the cell design. Specifically, this time is the time period from when precharging is complete to when the sense amplifier has reached 90% of its swing.

[0088] While the discharge times are asymmetric, the worst-case sensing times are on par with the RV cell with a conventional sense amplifier 130. Compared with the RV cell 10 with a conventional sense amplifier 130, the LE cell 40 is 10% slower. The effect on the total read time is an increase of just under 5%, however. The SE cell 70 is slightly faster not because the sense amplifier 140 is quicker, but because the bitline discharge time for the SE cell 70 is 50ps quicker than that of the RV cell 10, which is a by-product of the asymmetry of the SE cell 70. Furthermore, the RLE cell 120 has a worst-case sense time that is 2.5% slower than the RV cell 10, with the effect on total read time being near 1%. Interestingly, the HV cell 25 with a conventional sense amplifier 130 would be 26% slower.

[0089] An important side comment to be made is that the new sense amplifier 140 does not speed up the sensing for the RV 10 and HV 25 cells when compared to the sensing with the conventional sense amplifier. Indeed, the RV 10 and HV 25 cells with the new sense amplifier 140 have worst-case sense times that are 5% slower than the sense times with the conventional sense amplifier 130. Thus, in comparing the speed of the new cells with the new sense amplifier 140 to the conventional cells with the conventional sense amplifier 130, the comparison is fair and valid, because the new sense amplifier 140 on its own does not speed up the read access time of the conventional cells.

[0090] The LE 40 and SE 70 cells exhibit a write time increase of 19.4% and 25.3% respectively over the RV cell 10. The SLE 90 and SSE 100 cells exhibit an increase of 28.4% and 13.4% respectively, and the RLE 120 and RSE 110 exhibit an increase of 22.4% and 27.6% respectively. The increase in write times is of minor importance since the write times are all shorter than the read times of the associated cells and therefore the speed of the SRAM is dependent on the read time.

[0091] The present invention also analyzes two cache organizations that use asymmetric cell designs: statically biased and dynamic inversion. In the statically biased cache, the cells are simply replaced with asymmetric ones. This cache is statically biased to dissipate low leakage power only when it stores the preferred bit value '0'. What makes this cache successful is typical program behavior that exhibits a strong bias towards zero. Specifically, we observed that a level-1 data cache had an average 78.7% zeros in the data stream, and a level-1 instruction cache had an average of 62.9% zeros. Given this, the statically biased cache with the SE cells reduces leakage by 4.5X and 3.8X for an instruction and a data cache, respectively, compared to conventional symmetric cell

caches. The caches are 39Kbyte 4-way set associative caches. While programs with a higher fraction of '1's than '0's may exist, our SRAM would still dissipate much lower leakage power compared to the regular V_t cell cache.

[0092] In selective inversion, the values stored within a block can be inverted at a byte granularity (other granularities are possible). In this design, if a byte contains five or more ones it is inverted prior to storing it in the cache. This cache needs an additional inversion flag cell per byte that holds information on which bytes were inverted. Inversion happens at write time. Since stores are typically buffered in a write buffer and are only sent to the data cache on commit, there is plenty of time to decide and apply inversion if necessary. A logic flow diagram for this procedure is illustrated in FIGURE 15.

[0093] In yet another embodiment, an asymmetric SRAM cell design offers lower gate leakage with little or no impact on overall memory access time. In this asymmetric SRAM cell, an extra transistor is added to reduce the voltage across the gate of a leaky transistor to reduce leakage when the cell is storing a zero (the common case). As the gate oxide thickness gets thinner, gate tunneling leakage could surpass weak inversion and drain-induced-barrier-lowering leakage as the dominant leakage mechanism in future technologies.

[0094] Referring to FIGURE 16, there are multiple sources of gate leakage (see FIGURE 1) in an SRAM cell, but this embodiment aims to reduce the gate tunneling current through transistor N2. One reason for this is that 70% of cache bits are zeros, and thus it is more important to reduce leakage in one state. The cache bit binary values appear in parentheses throughout FIGURES 16-18. Second, the gate tunneling current through a PMOS transistor is an order of magnitude less than that through an NMOS. Also, edge directed tunneling (EDT) is an order of magnitude less than the on-gate leakage. These series of simplifications imply that the gate leakage through N2 is considered important.

[0095] This embodiment aims to reduce the gate leakage in SRAM cells. The previous asymmetric SRAM cells disclosed above reduce subthreshold leakage. The two approaches can also be combined in an orthogonal fashion with the dual- V_t cells to reduce both gate and subthreshold leakage.

[0096] Since gate leakage is exponentially related to VGS and VGD, one way to reduce gate leakage is to reduce the voltage on the storage nodes. A reduced VDD lowers the voltage at the storage nodes and thus reduces the gate leakage, as well as threshold leakage, in the cell. This technique, however, lowers the stability of the cell, and increases the delay and dynamic power consumption of the cell since VDD, must be switched to its nominal value before a read.

[0097] Another possibility is to slightly decouple the storage node from the gate of the pull-down transistor, so that the voltage across N2 can be lowered without reducing the voltage at the storage nodes. This is illustrated in FIGURE 16. In this pass cell (PC) 160, an NMOS pass transistor 162 has been inserted between the right storage node and the gate of N2.

[0098] When the cell is holding a '0', which is the common case, the N5 pass transistor 162 enters cutoff when its source voltage (gate of N2) is $V_{DD} - V_t$. Thus, the voltage across the gate of N2 has been reduced from VDD to $V_{DD} - V_t$, resulting in reduced direct tunneling leakage. This is

illustrated in FIGURE 17. Notice that the voltage at the storage nodes is not affected by this design. However, transistor N2's conductance has been reduced, which may affect the performance of the cell, and will be discussed below.

[0099] When the cell is holding a '1', as shown in FIGURE 18, the N5 pass transistor 162 plays no role in the functioning of the cell as an NMOS N5 pass transistor 162 is a good conductor of a logic '0'. There is, however, another source of direct tunneling leakage in this state. N5's (162) gate is at VDD, but it's source and drain are both and ground, thus there is direct tunneling leakage through N5's (162) gate. Since there are many more '0's than '1's, there is still a net leakage reduction.

[00100] At high values of t_{ox} , subthreshold leakage dominates, and there is no discernible decrease in total leakage. As t_{ox} is lowered to 1.4nm, the total leakage when storing a '0' has been reduced by 30% of the conventional cell's total leakage, but the total leakage when holding a '1' has been increased by 39%. Since 70% of SRAM cells are storing a '0', the total leakage of the array is reduced by 9%. By a t_{ox} of 1.3nm the total leakage when holding a '0' is reduced by 40%, while the total leakage when holding a '1' has increased by 50% leading to a total reduction of leakage by 13%.

[00101] When combined with VPL control, the total leakage savings become 30% at 1.3nm and when combined with a dual-Vt scheme, these savings become 40%. A 40% leakage reduction for a cache represents significant improvement.

[00102] If the cell was operating at lower temperatures, the total leakage savings would be even larger, since the subthreshold leakage would not be a large component of the leakage. For example, with the cell operating at 27°C at 1.3nm, the total leakage of the cache would be reduced by 37%.

[00103] The read access time of the pass cell is, however, degraded. When the cell is holding a '0', the bitline discharge along BL takes longer due to N2's lower conductance. The discharge time, which is only a small part of the total read access time, along BL is only 11% longer when the t_{ox} is 1.4nm. When the cell is holding a '1', there is no speed degradation along the BLB discharge path. There is actually a 4% speedup in the BLB discharge path due to the asymmetry in the cell.

[00104] The asymmetry in the discharge times can be used to have fast access times regardless of the value being stored. By using a new sense amplifier and a set of dummy bitlines, the read access times of the slow side of asymmetrical cells can be made to match the faster read time. To obtain fast read times irrespective of the data, a new sense amplifier was designed in. The new circuit uses a set of dummy bitlines, D and DB, which are connected to a column of cells that all hold a '1'. Thus D and DB are always fast and are used to trigger the reading of a logical '0' thus achieving fast access times when the slow bitline is discharging.

[00105] With the added pass-transistor in the pass cell, the time to flip the cell has slightly increased. This is due to the increased delay through N5 162, to turn off N2 when a '1' is being written into the cell. In the worst case, at large t_{ox} 's, there is a 36% increase in the flip time.

[00106] Another consideration with the cell design is its stability. There are two interrelated issues: read stability and noise margins. Intuitively, read stability indicates how likely it is to invert the cell's stored value when accessing it, and was computed as the ratio of I_{trip}/I_{read} , where I_{trip} is the current through the pull-down NMOS when the state of the cell is being reversed by injecting an

external current I_{test} and where I_{read} is the maximum current through the pass transistor during a read. The static noise margin (SNM) of an SRAM cell is defined as the minimum DC noise voltage necessary to flip the state of the cell. In our case, the stability of all cells was measured by simulation via both the Static Noise Margin (SNM) and the I_{trip}/I_{read} methods. Under both stability tests, the stability was first measured under nominal conditions, assuming no process variations.

[00107] Then, to measure stability under process variations, Monte-Carlo analysis was performed to obtain a distribution for the SNM and I_{trip}/I_{read} . For each cell, 500 scenarios for V_t and length were randomly generated, consistent with their joint distributions, and simulated. The mean and variance of the distribution were then estimated.

[00108] The SNM of the pass cell remains largely unchanged since it is a DC characteristic. By a tox of 1.4nm the SNM has only decreased by 0.25%, and by a tox 1.05nm, there is only a 4% drop.

[00109] The I_{trip}/I_{read} measure is actually improved by the inclusion of N5 162. Even at a tox of 1.3nm the pass cell shows a minimum of a 7% improvement in I_{trip}/I_{read} over the conventional cell. On the fast side of the cell, which has the limiting I_{trip}/I_{read} value, I_{trip} is increased due to the larger capacitance at the storage node. On the slow side of the cell is I_{trip} increased due to the slight decoupling of the positive feedback due to N5 162.

[00110] With lower values of tox however, the I_{trip}/I_{read} of the slow side of the cell becomes the limiting value. This is because the voltage on the gate of N2 is reduced due to the increased direct tunneling leakage, which in turn reduces the current sinking capability of N2, thus reducing I_{trip} .

[00111] The reason for designing the PC was to reduce the gate tunneling leakage, and this was accomplished by using a V_t drop from VDD to reduce the direct tunneling leakage. Thus if N5 162 had higher V_t , then the direct tunneling leakage savings will be increased. This technique would only reduce leakage in the '0' state. If on the other hand the voltage on PL is lowered, leakage in both the '0' and '1' state will be reduced.

[00112] Not only does the '0' leakage decrease because of the lower voltage on the gate of N2, but the '1' leakage decreases due to the reduced VGS and VGD on N5 162. With a 0.2V decrease in VPL, the total leakage when holding a '0' is reduced by an additional 10%, to 49% of the total leakage of a conventional cell. The total leakage of a cache would be reduced by an additional 17% to 70% of a conventional cache's leakage.

[00113] A decreased VPL leads to a lower voltage on the gate of N2, and thus the discharge time on BL increases. For example, with VPL dropping to IV, the discharge time is 30% longer than the CC. The increase in discharge time is, however, inconsequential because of the special sense-amplifier, which matches the time of the slow side to the fast side of the cell, whose performance does not change.

[00114] When decreasing VPL, the flip times of the pass cell actually decrease. This is due to the reduced feedback within the cell. For example, when the cell is holding a '0' and a '1' is being written into it, N2 cannot sink as much current since its gate voltage is not at VDD, and thus the write occurs quicker.

[00115] As VPL decreases, and the voltage on the gate of N2 decreases, N2's conductance drops which will impact the stability of the pass cell. Even at a VPL of 1 V there is only a 13% drop in SNM.

[00116] For the Itrip/Iread measure, as VPL is decreased, the 'trip on the fast side of the cell increases because of the extra delay through N5, thus increasing Itrip/Iread. On the slow side of the cell, the limiting side, 'trip decreases due to the reduced conductance of N2, thus lowering Itrip/Iread. Thus, the different sides of the cell experience different stability characteristics.

[00117] By decreasing VPL we are trading the increased stability of the pass cell for reduced leakage. With a 0.2V drop on VPL, the total leakage of the cache reduces from 87% of a conventional cache to 70% of the conventional cache with little change in stability compared to that of a conventional cell.

[00118] As tox becomes thinner, subthreshold leakage will compose a smaller, yet still important part of total static power consumption. As mentioned earlier, the pass cell can be combined with the Stability-Speed Enhanced (SSE) and Resized Leakage Enhanced (RLE) asymmetric dual-Vt cells described above to obtain the PC-SSE and PC-RLE cells to lower both gate and subthreshold leakage. Not only are the leakage savings orthogonal, but so are the performance and stability of the dual-Vt pass-cell.

[00119] At high tox 's where subthreshold leakage dominates, the PC-SSE and PC-RLE cells have the same total leakage as the SSE and RLE cells, a 2X and 17X reduction in total leakage. As gate leakage becomes a more important part of the total leakage, the PC-SSE and PC-RLE cells have better leakage than both the dual-Vt cells that they derived from, and from the single-Vt PC-LV cell. At a tox of 1.3nm the PC-SSE and PC-RLE save an additional 10% of total leakage compared to the pass cell.

[00120] Since the pass transistor is on the slow side of the cell, the discharge time on the east side of the cell is unaffected when transforming the SSE to the PC-SSE and when changing the RLE to the PC-RLE. Thus due to the sense amplifier presented in, there is no added speed degradation. The PC-LV cell, however, had a slight speedup in the BILE discharge time due to asymmetry inherent in the cell. This slight speedup remains in the PC-SSE cell which is 4% faster than the CC. The RLE cell is 4% slower than the CC cell, but due to the added asymmetry of N5, the PC-RLE is now 1% faster.

[00121] By combining the dual-Vt cells and the N5 pass-transistor 162 the flip times for the cells increase. The flip time increase is nearly equal to the separate flip time increases associated with the dual-Vt and pass cell design. Regardless, the flip time increases are only on the order of tens of picoseconds.

[00122] The increased stability in the SSE and RLE cells help to increase the stability of the pass cell when it is transformed to the PC-SSE and PC-RLE cell. At a tox of 1.3nm the Itrip/Iread of the dual-Vt cells is at least 12% better and the SNM of the PC cell increases; the PC-SSE cell only has a 7% decrease in SNM, while the PC-RLE cell has a 12% increase.

[00123] Thus, the pass cell reduces direct tunneling leakage through one of the pull-down NMOS transistors because of three key observations. First, gate leakage through a PMOS is an order of magnitude less than through an NMOS, EDT tunneling is an order of magnitude less than the direct

tunneling leakage, and cache-resident memory values of ordinary programs exhibit a strong bias towards zero at the bit level.

[00124] The pass cell can be combined with dual-Vt design to reduce both subthreshold and gate leakage. There are multiple design possibilities that have different performance/leakage/ stability characteristics. At a tox of 1.3nm, the best design reduces leakage to 60% of that of a conventional cell with no performance degradation and comparable stability. There is, however, a 16.6% increase in cell area.

[00125] Furthermore, the leakage savings with this design are orthogonal to leakage savings incurred by turning off parts of the cache or reducing VDD on the cell supply, or allowing the bitline voltage to float.

[00126] The present invention presents a novel approach that combines both circuit and architecture level techniques for drastically reducing leakage power dissipation. A key observation behind the present invention is that cache-resident memory values of ordinary programs exhibit a strong bias towards zero or one at the bit level. The present invention has introduced a family of high-speed asymmetric dual-Vt SRAM cell designs that exploit this bit-level bias to reduce leakage power while maintaining high performance.

[00127] Various asymmetric cells offer different performance/leakage/stability characteristics. The SE cell reduces leakage power by at least 2X and by 7X in the preferred state. It is as fast as the conventional, RV, SRAM cell. By comparison, the LE cell reduces leakage by at least 7X and by about 70X in the preferred state. Its total read time is only 5% higher than the SE and RV cells. These latter two cells have lower stability than LE under both the SNM and the Itrip/Iread tests. Four other cells that compensate for stability were also designed, two by choosing different combinations of threshold voltages for the cell transistors, and two by changing some transistor sizes. The SSE cell reduces leakage power by 1.9X and 2.3X in the preferred state with no performance degradation, and the SLE cell reduces leakage power by 2.3X and 7X in the preferred state with only a 5% increase in read access times. The SSE and SLE cells have comparable stability to the RV cell. The RLE cell reduces leakage by 58X in the preferred state and by 7X in the other state with only a 1% increase in read access time, and an area increase of about 2.4%. The RSE cell reduces leakage by about 7X in the preferred state, and 2X in the other state. It has no performance degradation, but has an area increase of about 2.9%. The RLE and RSE cells have comparable stability to the RV cell. By comparison, an all high Vt cell reduces leakage power by about 70X while its bitline discharge time is 60% slower than the SE and RV cells.

[00128] The present invention also presents two cache organizations that use either a static bias towards zero, or dynamic, selective inversion to maximize the number of cache bits that are zero. While the reduction possible with either technique depends on application behavior, the statically biased cache with the SE cells reduces leakage by 4.5X and 3.8X for an instruction and a data cache, respectively, as compared to conventional symmetric-cell caches.

[00129] The present invention further presents an SRAM cell designed to lower gate tunneling leakage while maintaining high performance and comparable noise margins and stability.

This design can even be used in an orthogonal fashion with the dual-Vt cells described above to increase the leakage savings and stability.

[00130] The preceding description has focused on SRAM cell designs that were comprised of six transistors. The principles of the present invention were described and applied to a six transistor design for ease of illustration. It should be noted, however, that the same asymmetric principles of the present invention may also be applied to other SRAM cell designs including, but not limited to, those comprised of four transistors and two resistors.

[00131] It is the asymmetric nature of the present invention that provides the novelty and uniqueness rather than a particular SRAM architecture. Thus, SRAM cell designs, as well as sense amplifiers and SRAM devices comprised of arrays of SRAM cells, that exhibit asymmetric transistor design characteristics are considered within the scope of the present invention.

[00132] Specific embodiments of an invention are described herein. One of ordinary skill in the circuit design and computing arts will quickly recognize that the invention has other applications in other environments. In fact, many embodiments and implementations are possible. The following claims are in no way intended to limit the scope of the invention to the specific embodiments described above.

CLAIMS

We claim:

1. An asymmetric SRAM cell for storing a binary variable, the asymmetric SRAM cell having reduced leakage power with respect to a comparable symmetric SRAM cell when the asymmetric SRAM cell stores a binary variable representing a predetermined binary value, the asymmetric SRAM cell comprising:

a plurality of transistors operably coupled and configured as an asymmetric SRAM cell, wherein the plurality of transistors include at least one first type of transistor and at least one second type of transistor that is weaker than the first type of transistor, such that the configuration of the asymmetric SRAM cell achieves reduced leakage power with respect to a symmetric SRAM cell having the first type of transistor only.

2. The asymmetric SRAM cell of claim 1 wherein at least one of the second type of transistor is selected from among the group consisting of:

a transistor having a higher voltage threshold (V_t) as compared to the voltage threshold (V_t) of the first type of transistor;

a transistor having a decreased channel width as compared to the channel width of the first type of transistor; and

a transistor having an increased channel length as compared to the channel length of the first type of transistor.

3. A sense amplifier for coupling with an asymmetric SRAM cell that provides faster access times when the asymmetric SRAM cell stores a first predetermined binary value, said sense amplifier comprised of:

a first pair of cross coupled inverters across a bitline (BL) and a bitline bar (BLB);

a second pair of cross coupled inverters operably coupled with the first pair of cross coupled inverters;

a plurality of additional transistors forming a dummy column of cells that store a second predetermined binary value at all times wherein during a read operation of the SRAM cell one of the dummy cells will have its wordline asserted, said dummy column of cells operably coupled with the first pair of cross coupled inverters; and

four inputs operably coupled with a subset of transistors of the sense amplifier wherein the inputs include the BL, the BLB that derive from the SRAM cell, a dummy bit line (D), and a dummy bitline bar (DB) that are input to the dummy cells such that D is input to the sense amplifier on the same side as BLB while DB is input to the sense amplifier on the same side as BL.

4. The sense amplifier of claim 3 wherein at least one of the transistors coupled with BL and BLB have higher transconductance characteristics than at least one of the transistors coupled with D and DB.

5. The sense amplifier of claim 3 wherein at least one of the transistors coupled with BL and BLB are selected from among the group consisting of:

transistors having a lower voltage threshold (V_t) as compared to the voltage threshold (V_t) of the transistors coupled with D and DB;

transistors having a increased channel width as compared to the channel width of the transistors coupled with D and DB; and

transistors having a decreased channel length as compared to the channel length of the transistors coupled with D and DB.

6. An SRAM device comprising:

an array of SRAM cells wherein each SRAM cell stores a binary variable representing a predetermined binary value, and each SRAM cell is an asymmetric SRAM cell having reduced leakage power with respect to a comparable symmetric SRAM cell, each asymmetric SRAM cell comprising:

a plurality of transistors operably coupled and configured as an asymmetric SRAM cell, wherein the plurality of transistors include at least one of a first type of transistor and at least one of a second type of transistor that is weaker than the first type of transistor, such that the configuration of each asymmetric SRAM cell achieves reduced leakage power with respect to a symmetric SRAM cell having the first type of transistor only.

7. The SRAM device of claim 6 wherein the array of SRAM cells in the SRAM device comprises an SRAM device selected from the group consisting of a direct store SRAM device and a selectively inverted SRAM device.

8. The SRAM device of claim 6 wherein the array of SRAM cells in the SRAM device comprises a cache memory selected from the group consisting of a direct store cache memory and a selectively inverted cache memory.

9. A combination SRAM device and sense amplifier comprising:

an array of SRAM cells wherein each SRAM cell stores a binary variable representing a predetermined binary value, and wherein each SRAM cell is an asymmetric SRAM cell having reduced leakage power with respect to a comparable symmetric SRAM cell, each asymmetric SRAM cell comprising:

a plurality of transistors operably coupled and configured as an asymmetric SRAM cell, wherein the plurality of transistors include at least one of a first type of transistor and at least one of a second type of transistor that is weaker than the first type of transistor, such that the configuration of each asymmetric SRAM cell achieves reduced leakage power with respect to a symmetric SRAM cell having the first type of transistor only; and

at least one sense amplifier comprised of:

a first pair of cross coupled inverters across a bitline (BL) and a bitline bar (BLB);

a second pair of cross coupled inverters operably coupled with the first pair of cross coupled inverters;

a plurality of additional sense amplifier transistors forming a dummy column of cells that store a second predetermined binary value at all times wherein during a read operation of the SRAM cell one of the dummy cells will have its wordline asserted, said dummy column of cells operably coupled with the first pair of cross coupled inverters; and

four inputs operably coupled with a subset of the sense amplifier transistors wherein the inputs include the BL, the BLB that derive from the SRAM cell, a dummy bit line (D), and a dummy bitline bar (DB) that are input to the dummy cells such that D is input to the sense amplifier on the same side as BLB while DB is input to the sense amplifier on the same side as BL.

10. The combination SRAM device and sense amplifier of claim 9 wherein the sense amplifier transistors coupled with BL and BLB have higher transconductance characteristics than the sense amplifier transistors coupled with D and DB.

11. The combination SRAM device and sense amplifier of claim 9 wherein at least one of the sense amplifier transistors coupled with BL and BLB are selected from among the group consisting of:

transistors having a lower voltage threshold (V_t) as compared to the voltage threshold (V_t) of the transistors coupled with D and DB;

transistors having a increased channel width as compared to the channel width of the transistors coupled with D and DB; and

transistors having a decreased channel length as compared to the channel length of the transistors coupled with D and DB.

12. The combination SRAM device and sense amplifier of claim 9 wherein the SRAM device comprises an SRAM device selected from the group consisting of a direct store SRAM device and a selectively inverted SRAM device.

13. The combination SRAM device and sense amplifier of claim 12 wherein the array of SRAM cells in the SRAM device comprises a cache memory selected from the group consisting of a direct store cache memory and a selectively inverted cache memory.

14. An asymmetric static random access memory (SRAM) cell operable with a supply voltage to store a one or a zero, the asymmetric SRAM cell further comprising:

at least two cross-coupled inverters, one having an output electrically connected to a bit line bar when a word line is held high and the other having an output electrically connected to a bit line when a word line is held high and further comprising a pull-down transistor; and

a pass transistor connected to a gate of the pull-down transistor so that a voltage across the gate is reduced relative to the supply voltage, thereby reducing leakage through the gate when the asymmetric SRAM cell is storing a zero.

15. The asymmetric SRAM cell of claim 14 wherein a voltage at a gate of the pass transistor is reduced relative to the supply voltage to further reduce leakage through the gate of the pull-down transistor.

16. The asymmetric SRAM cell of claim 14 wherein the first and second inverters comprise a plurality of transistors further comprising at least one first type of transistor and at least one second type of transistor that is weaker than the first type of transistor.

17. The asymmetric SRAM cell of claim 16 interconnected with a plurality of like SRAM cells and a sense amplifier further comprising pairs of cross-coupled inverters and a plurality of sense amplifier transistors forming a dummy column of cells.

ABSTRACT

[00133] Asymmetric SRAM cell designs exploiting data storage patterns found in ordinary software programs wherein most of the bits stored are zeroes for data and instruction streams. The asymmetric SRAM cell designs offer lower leakage power with little impact on latency. In asymmetric SRAM cells, selected transistors are "weakened" to reduce leakage current when the cell is storing a zero. Transistor weakening may be achieved by using higher voltage threshold transistors, by varying transistor geometries, or other means. In addition, a novel sense amplifier design is provided that leverages the asymmetric nature of the asymmetric SRAM cells to offer cell read times that are comparable with conventional symmetric SRAM cells. Lastly, cache memory designs are provided that are based on asymmetric SRAM cells offering leakage power reduction while maintaining high performance, comparable noise margins, and stability with respect to conventional cache memories.

- Subthreshold Leakage
..... On Gate Direct Tunneling
→ Edge Directed Tunneling

FIG. 1







FIG. 2









FIG. 7




FIG. 8




FIG. 9







FIG. 10





130

FIG. 13

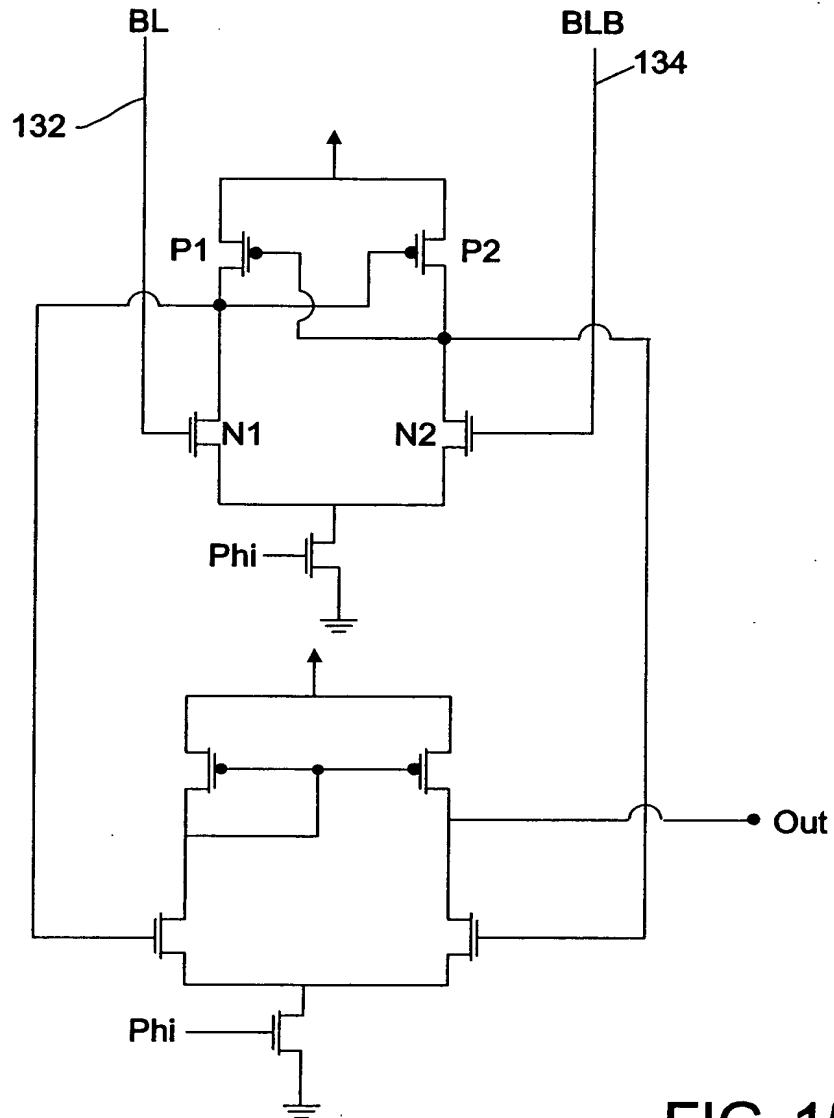




FIG. 15



140



160




FIG. 18

